

Computing rational decisions in extensive games with limited foresight

Paolo Turrini

Department of Computing, Imperial College London

May 31, 2016

Abstract

We introduce a class of extensive form games where players might not be able to foresee the possible consequences of their decisions and form a model of their opponents which they exploit to achieve a more profitable outcome. We improve upon existing models of games with limited foresight, endowing players with the ability of higher-order reasoning and proposing a novel solution concept to address intuitions coming from real game play. We analyse the resulting equilibria, devising an effective procedure to compute them.

1 Introduction

While game theory is a predominant paradigm in Artificial Intelligence, the tools it provides to analyse real game play still abstract away from many essential features. One of them is the fact that in a wide range of extensive games of perfect information (e.g., Chess), humans (and supercomputers) are generally not able to fully assess the consequence of their own decisions and need to resort to a judgment call before making a move. As acclaimed game theorist Ariel Rubinstein puts it, "modeling games with limited foresight remains a great challenge" and the game-theoretic frameworks developed thus far "fall short of capturing the spirit of limited-foresight reasoning" (Rubinstein, 2004, p.134).

On the contrary, the AI approach to game-playing builds upon the assumption that complex extensive games like Chess or Go are *theoretically* games of perfect information, but this is only marginally relevant for practical purposes, and the backwards induction solution is of little help in predicting how such games are actually played in practice - a point also raised in Joseph Halpern's AAMAS 2011 invited talk "Beyond Nash Equilibrium: Solution Concepts for the 21st Century" (Halpern, 2008). Decisions are instead taken using heuristic search (e.g., monte-carlo tree search) under various constraints, such as time or memory (Russell and Wefald, 1991) (Russel and Norvig, 2012).

The problem. Search methods are a framework to handle limited foresight and are widely used for decision-making in real game-play, but a game-theoretic analysis of their equilibrium behaviour is still missing. In particular, we lack the tools to analyse what will happen in complex extensive games of perfect information where players are not able to resort to backwards induction reasoning but to possibly faulty and incomplete heuristic. What is more, the enormous effort to construct players with "opponent modelling" in the AI community (e.g., (Schadd *et al.*, 2007) (Donkers *et al.*, 2001)) still lacks solid game-theoretic foundations.

Our contribution. We introduce games in which players might not be able to foresee the consequence of their strategic decisions all the way up to the terminal nodes and evaluate intermediate nodes according to a concrete heuristic search method. On top of that they can reason about other players' limited foresight and evaluation criteria: they are endowed with higher-order beliefs about what their opponents can perceive of the game and how they evaluate it, beliefs about what their opponents believe the others can see and how they evaluate it, and so forth. To analyse these games, we propose a new solution concept which combines higher-order reasoning about players' limited foresight and evaluation criteria. *The guiding principle for players' behaviour is that each of them chooses a strategy in the game she sees that is a best response to the belief about what the other players can see and how they evaluate it.* We show constructively (Algorithms 1-4) that this solution concept always exists (Theorem 1) and is a strict generalization of other known ones, e.g., backwards induction. As we will observe, the unbounded chain of beliefs underlying our rationality constraints can be finitely represented and - rather surprisingly - effectively resolved (Proposition 13).

Related literature. In recent years an innovative tradition has emerged in game theory, aiming at capturing situations in which players are unaware of parts of the game they are playing and might even think to be playing a different game from the real one. Halpern and Rêgo (Halpern and Rêgo, 2006), for instance, study models of unawareness of elements of the game played (e.g., other players). Yossi Feinberg (Feinberg, 2012) approaches similar problems from a syntactic perspective. Simultaneously, the interplay between belief and awareness in interactive situations is analysed in a series of papers by Heifetz, Meier and Schipper (Heifetz *et al.*, 2006), (Heifetz *et al.*, 2013a), (Heifetz *et al.*, 2013b).

It should be noted that even though all these frameworks abstractly allow to talk about unawareness of *some* terminal histories in a game, none of them comes equipped with a solution concept capturing limited foresight reasoning.

A framework that comes closest, perhaps, to this is *Games with Short Sight* (Grossi and Turrini, 2012), a well-behaved collection of games with awareness, in which players of an extensive game make choices without knowing the consequences of their actions and base their decisions on a (possibly incorrect) evaluation of intermediate game positions.

Games with Short Sight (GSSs) have been studied in relation with a solution concept called *sight-compatible backwards induction*: as players might not be able to calculate all possible moves up to the terminal nodes, they play rationally in a *local* sense, executing moves that are backwards induction moves in their own sight, therefore *safely* assuming their opponents see as much of the game as they do.

However, sight-compatible backwards induction precludes any sort of *opponent modelling*, as players are not allowed to have a non-trivial belief about what their opponents perceive. Thus, the tools developed in (Grossi and Turrini, 2012) to analyse GSSs only allow players to play approximately or inaccurately, they don't allow players to exploit their opponents' believed weaknesses. Besides, GSSs employ heuristics which are not grounded in practical game-play. Essentially, players come equipped with a preference relation over all histories of the game.

We will avoid strong rationality requirements of this kind, by introducing a significantly higher level of complexity in players' reasoning - notably their ability of forming an "opponent model" - which, it turns out, still remains computationally manageable. Also players' preference relations will not be taken as given, but derived from concrete search methods.

An important research line in AI that has similarities with our approach is interactive POMDPs (Gmytrasiewicz and Doshi, 2005), which is able to incorporate higher-order

epistemic notions in multi-agent decision making, with focus on learning and value/policy iteration. These graph-like models are generally highly complex - in fact the whole approach is known to suffer from severe complexity problems when it comes to equilibrium analysis and approximation methods have been devised to (partially) address them (Doshi and Gmytrasiewicz, 2009),(Sonu and Doshi, 2015). Instead, we present a full-blown game-theoretic model of limited foresight that allows for higher-order epistemic notions and yet keeps equilibrium computation within polynomial time.

Paper Structure. Section "Games with limited foresight" recalls useful formal notation and definitions from the literature upon which we build and introduces the mathematical structures we will be working on, Monte-Carlo Tree Games. Section "Rational beliefs and limited foresight" studies the higher-order extension thereof, Epistemic Monte-Carlo Tree Games. Specifically, we go on and define a new solution concept which takes this higher-order dimension into account and we then show the existence of the new equilibria through an efficient (P-TIME) algorithm. Section "Conclusion and potential developments" summarises our findings and hints at new research avenues opening up in our framework.

2 Games with limited foresight

We start out with the definition of extensive games, on top of which we build the models of limited foresight.

Extensive Games An **extensive game form** (Osborne and Rubinstein, 1994) is a tuple

$$(N, H, t, \Sigma_i, o)$$

where [1] N is a finite non-empty set of **players**. [2] H is a non-empty prefix-closed set of sequences, called **histories**, drawn from a set A of **actions**. A history $(a^k)_{k=1,\dots,K} \in H$ is called **terminal** history if it is infinite or if there is no a^{K+1} such that $(a^k)_{k=1,\dots,K+1} \in H$. The set of terminal histories is denoted Z . A history h is instead called **quasi-terminal** if for each $a \in A$, if $(h, a) \in H$, then (h, a) is terminal. If $h \in H$ is a prefix (resp., strict prefix) of $h' \in H$ we write $h \trianglelefteq h'$ (resp., $h \triangleleft h'$). With $A_h = \{a \in A \mid (h, a) \in H\}$ we denote the set of actions following the history h . The restriction of $H' \subseteq H$ to $h \in H$, i.e., $\{(h, h') \in H \mid (h, h') \in H'\}$ is denoted $H'|_h$. [3] $t : H \setminus Z \rightarrow N$ is a **turn function**, which assigns a player to each non-terminal history, i.e., the player who *moves* at that history. [4] Σ_i is a non-empty set of **strategies**. A strategy of player i is function $\sigma_i : \{h \in H \setminus Z \mid t(h) = i\} \rightarrow A_h$, which assigns an action in A_h to each non-terminal history for which $t(h) = i$. [5] o is the **outcome function**. For each strategy profile $\sigma = \prod_{i \in N} (\sigma_i)$, the outcome $o(\sigma)$ of Σ is the terminal history that results when each player i follows the precepts of σ_i .

An **extensive game** is a tuple $\mathcal{E} = (\mathcal{G}, \{u_i\}_{i \in N})$, where \mathcal{G} is an extensive game form, and $u_i : Z \rightarrow \mathbb{R}$ is a utility function for each player i , mapping terminal histories to reals. We denote $\succeq_i \subseteq Z \times Z$ the induced total preorder over Z and $\text{BI}(\mathcal{E})$ the set of backwards induction histories of extensive game \mathcal{E} , computed with the standard procedure (Osborne and Rubinstein, 1994, Proposition 99.2).

Sight Functions and Forked Extensions On top of the extensive game structure, each player moving at the certain point in the game is endowed with a set of histories that he or she can see from then on.

Consider an extensive game $\mathcal{E} = (\mathcal{G}, \{u_i\}_{i \in N})$. A (**short**) **sight function** for \mathcal{E} (Grossi and Turrini, 2012) is a function

$$s : H \setminus Z \rightarrow 2^H \setminus \emptyset$$

associating to each non-terminal history h a finite non-empty and prefix-closed subset of all the histories extending h , i.e., histories of the form (h, h') . We denote $H \lceil_h = s(h)$ the **sight restriction** on H induced by s at h , i.e., the set of histories in player $t(h)$'s sight, and $Z \lceil_h$ their terminal ones. Intuitively, the sight function associates any choice point with those histories that the player playing at that choice point *actively* explores.

In (Grossi and Turrini, 2012) the problem of evaluating intermediate positions is resolved by assuming the existence of an arbitrary preference relation over these nodes, which is common knowledge among the players. What we do instead is to introduce an extension of sight functions that models the evaluation obtained by a concrete search procedure. The idea is that in order to evaluate intermediate positions, each player carries out a selection and a random exploration of their continuations, all the way up to the terminal nodes. The information obtained is used as an estimate of the value of those positions. This is an encoding of a basic Monte-Carlo Tree Search (Browne *et al.*, 2012).

Let (\mathcal{E}, s) be a tuple made by an extensive game \mathcal{E} and a sight function s . Sight function s^* is called a **forked extension** of sight function s if the following holds:

- $s(h) \subseteq s^*(h)$ i.e., the forked extension prolongs histories in the sight it extends;
- For \lceil^* being the sight restriction calculated using s^* as sight function, we have that: if $h \in s^*(h) \setminus Z$ then there exists $h' \in s^*(h)$ such that $h \triangleleft h'$, i.e., $s^*(h)$ is made of histories that go all the way up to the terminal nodes.¹

A **Monte-Carlo Tree Game** (MTG) is a tuple $S = (\mathcal{E}, s, s^*)$ where $\mathcal{E} = (\mathcal{G}, \{u_i\}_{i \in N})$ is an extensive game, s a sight function for \mathcal{E} and s^* a forked extension of s . We denote $S \lceil_h = (\mathcal{G} \lceil_h, \{u_i \lceil_h\}_{i \in N})$ the sight restriction of S induced by s at h , where $\mathcal{G} \lceil_h$ is the game form \mathcal{G} restricted to $H \lceil_h$ and the utility function $u \lceil_h : N \times Z \lceil_h \rightarrow \mathbb{R}$ is constructed as follows. For each $i \in N, g \in Z \lceil_h$, we have:

$$u_i \lceil_h(g) = \text{avg}_{z \in Z \lceil_h^*, g \triangleleft z} u_i(z)$$

So the utility function at terminal histories in a sight is computed by taking the average² of the histories contained in its forked extension. Notice the following important point: histories in the forked extension are truly treated as "random" explorations, with no rationality assumptions whatsoever, in order to construct a preference relation over $Z \lceil_h$. Sight-restriction is applied to players, turn function, strategies and outcome function in the obvious way. Summing up, each structure $(\mathcal{G} \lceil_h, \{u_i \lceil_h\}_{i \in N})$ is an extensive game, intuitively the part of the game that the player moving at h is able to see, where the terminal histories are evaluated with a monte-carlo heuristic.

¹A further natural constraint on forked sight functions is that of monotonicity, i.e., players do not forget what they have calculated in the past. Formally s is *monotonic* if, for each h, h' such that $t(h) = t(h')$ and $h \triangleleft h'$, we have that $s^*(h)|_{h'} \subseteq s^*(h')$. Albeit natural, this assumption is not needed to prove our results.

²Averaging has the sole purpose of simplifying notation and analysis, which carries over to *any* aggregator, with or without lotteries. Besides, it comes along with a few desirable properties, notably the fact that forked extensions *never miss dominated continuations*, i.e., moves that ensure a gain no matter what the opponents do. For quantified restrictions on aggregators cfr. for instance (van Benthem *et al.*, 2011).

The solution concept proposed in (Grossi and Turrini, 2012) to analyse GSSs is *sight-compatible backwards induction*: a choice of strategy, one per player, that is consistent with the subgame perfect equilibrium of each sight-restricted game. We can encode it as follows.

Definition 1 (Sight-compatible BI) Let S be a MTG. A strategy profile σ is a *sight-compatible backwards induction* if at each $h \in H$, there exists a terminal history $z \in Z \upharpoonright_h$ such that $h\sigma(h) \trianglelefteq z$ and $z \in \mathbf{BI}(S \upharpoonright_h)$. The set of sight-compatible backwards induction outcomes of S is denoted $\mathbf{SCBI}(S) \subseteq Z$.

Thus, a sight compatible backwards induction is a strategy profile σ that, at each history h , recommends an action a that is among the actions initiating a backwards induction history within the sight of the player moving at h . This, notice, is different from the backwards induction solution of the whole game, because players evaluation of intermediate nodes might not be a *correct assessment* of the real outcomes of the game. Grossi and Turrini show that the SCBI solution always exists, even in infinite games.

Despite their effort in modelling more *procedural* aspects of game play, though, GSSs still lack non-trivial opponent modelling, i.e., players allowing for their opponents to “miss” future game developments and evaluate game positions differently (or any higher-order iteration of this belief), while adjusting their behaviour accordingly.

The rest of the paper is devoted to extending MTGs with more realistic but highly more complex reasoning patterns, generalising both GSSs and SCBI. This, it turns out, does not prevent us from having appropriate well-behaved solution concepts which generalise classical ones, such as backwards induction.

3 Rational beliefs and limited foresight

We now introduce an extension of MTGs, where players are allowed for the possibility of higher-order opponent-modelling, i.e., to have an explicit belief about what other players can see and how they evaluate it, a belief about what other players believe other players can see and how they evaluate and so forth, *compatibly with players’ sight*. We study a solution concept for these games and relate it to known ones from the literature.

3.1 Players’ sights and belief chains

Let us introduce the idea behind higher-order opponent modelling in MTGs using an example. We will then move on to define the notions formally.

Example 2 (An intuitive solution) Consider the game shown in Figure 1. Three players, Ann, Bob and Charles, move at histories marked A , B and C , respectively. The circle surrounding history A indicates what Ann believes she can see from history A , which we write $b(A)$. This, intuitively, coincides what Ann can actually see, i.e., it equals $s(A)$, Ann’s sight at history A . What should Ann do in this situation? This depends on what Ann believes will happen next. If Ann knew this, her choice would only be a maximization problem: finding the action that, given what will happen in the future, gets her the maximal outcome, according to her evaluation from A - which we write \succsim_{Ann}^A . To find out what Charles will do, Ann considers her belief about what Charles can see from C , which we indicate with $b(A)b(C)$. Note this may have nothing to do with what Charles actually sees from C , i.e., $b(C)$. In Figure 1, for instance, Ann believes that Charles can only see d from C . The question of what Charles will

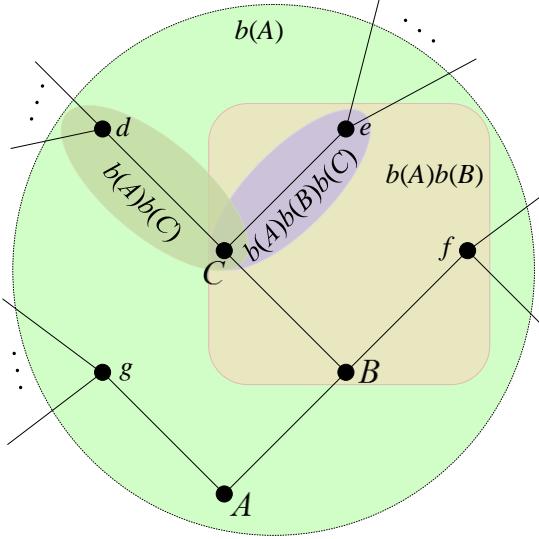


Figure 1: A belief structure (modulo evaluations)

do is then easily answered, even without considering his preference relation \succ_{Ann}^{AC} , what Ann believes Charles wants from history C . Charles, according to Ann, will certainly go to d . The next question is: what will Bob do? This, again, will depend on $b(A)b(B)$, the portion of Ann's sight that Ann believes Bob can see from B and on \succ_{Ann}^{AB} , the preferences Ann believes Bob has at B . But, at least according to Ann, Bob can also see that Charles can make moves. So, for Bob to decide what to do, he must first find out what Charles will do - $b(A)b(B)b(C)$ - according to \succ_{Ann}^{ABC} . This is also an easy task, since e is the only option. The choice at $b(A)b(B)b(C)$ is then determined, but so is then the choice at $b(A)b(B)$. Now all that is left for Ann to do is to solve her maximization problem, determining the choice at $b(A)$.

Now we concentrate on turning the intuitions in the example into formal definitions. To do so, we introduce the notion of *history-sequence*. A history-sequence is a formal device that allows to represent higher-order beliefs about other opponents, consistently with a players' sight.

Definition 3 (History-Sequences) Consider a MTG $S = ((N, H, t, \Sigma_i, o), \{u_i\}_{i \in N}, s, s^*)$. A history-sequence q of S is a sequence of histories of the form $(h_0, h_1, h_2, \dots, h_k)$ such that

- $h_j \in H|_{h_0}$ for every $j \in \{1, 2, \dots, k\}$, i.e., histories following h_0 in the sequence are histories within the sight of the player moving at h_0 ;
- $h_j \triangleleft h_{j+1}$ for each j with $0 \leq j < k$, i.e., each history is a strict postfix of the ones with lower index;

The underlying idea behind this definition is to consider the higher-order point of view of the player moving at h_0 . Expressions of the form $(h_0, h_1, h_2, \dots, h_k)$ encode the belief that player moving at h_0 holds about the belief that player moving at h_1 holds about the belief that player moving at h_2 holds ... about what the player moving at h_k can see and

what the evaluation is of the corresponding terminal histories. We use \mathbf{Q} to denote the set of history-sequences of S .

Building upon the notion of history-sequence, we can define what we call *sight-compatible* belief structures, associating each history-sequence with a set of histories and an evaluation over the terminal ones in this set.

Definition 4 (Sight-compatible belief structures) Let S be a MTG. A *sight-compatible belief structure* \mathbf{B} for S is a tuple $(\mathbf{B}_H, \mathbf{B}_P)$ such that \mathbf{B}_H is a function $\mathbf{B}_H : \mathbf{Q} \rightarrow 2^H$, associating to each history-sequence $(h_0, h_1, h_2, \dots, h_k)$ a set of histories in $s(h_0)$ extending h_k , and \mathbf{B}_P is a function $\mathbf{B}_P : \mathbf{Q} \rightarrow 2^Z$ associating to each history-sequence \mathbf{q} a set of terminal histories extending histories in $\mathbf{B}_H(\mathbf{q})$. \mathbf{B} satisfies the following conditions:

- (Corr) $\forall \mathbf{q} \in \mathbf{Q}$ with $\mathbf{q} = (h_0)$, then $\mathbf{B}_H(\mathbf{q}) = H|_{h_0}$ whenever $t(h_k)$ the belief of a player about what he himself can see is correct.³
- (Mon of \mathbf{B}_H) $\forall \mathbf{q}, \mathbf{q}' \in \mathbf{Q}$, if $\exists h' \in \mathbf{B}_H(\mathbf{q})$ s.t., $\mathbf{q}' = (\mathbf{q}, h')$, then $\mathbf{B}_H(\mathbf{q}') \subseteq \mathbf{B}_H(\mathbf{q})|_{h'}$, i.e., if a player believes someone is able to perceive a portion of the game, then he is able to perceive that portion himself.
- (Mon of \mathbf{B}_P) $\forall \mathbf{q}, \mathbf{q}' \in \mathbf{Q}$, if $\exists h' \in \mathbf{B}_H(\mathbf{q})$ s.t., $\mathbf{q}' = (\mathbf{q}, h')$, then $\mathbf{B}_P(\mathbf{q}') \subseteq \mathbf{B}_P(\mathbf{q})|_{h'}$, i.e., if a player believes someone is able to explore a position, then he is able to perceive that exploration himself.

For $\mathbf{q} = (h_0, h_1, h_2, \dots, h_k)$, $\mathbf{B}_P(\mathbf{q})$ denotes the higher-order beliefs (in the order given by \mathbf{q}) about how player moving at h_k is evaluating the terminal histories in $\mathbf{B}_H(\mathbf{q})$ under the unique forked extension of s whose terminal histories are $\mathbf{B}_P(\mathbf{q})$. $\succeq_i^{\mathbf{B}_P(\mathbf{q})}$ denotes the induced preference relation, one per player.

The conditions above, we argue, are most natural constraints on sight-compatible higher-order beliefs. For the time being we do not commit ourselves to any other constraints on either \mathbf{B}_P or \mathbf{B}_H , but we acknowledge that different contexts may warrant further constraints on both.

Definition 5 (Epistemic Monte-Carlo Tree Games) An Epistemic Monte-Carlo Tree Games (EMTGs) is a tuple $\mathcal{S} = (S, \mathbf{B})$ where S is a MTG and \mathbf{B} a sight-compatible belief structure for S .

An EMTG is obtained by assigning a sight-compatible belief structure to a MTG. One should observe how sight-compatible belief structures induce, at each history, a whole collection of extensive games, one for each possible history-sequence. For instance, the one resulting from Ann's sight and her evaluation, the one resulting from Ann's belief about Bob's sight and his evaluation and so forth. Structures of the form $\mathcal{S}|_{\mathbf{B}(\mathbf{q})}$ can now be naturally defined, as restrictions induced by $\mathbf{B}(\mathbf{q})$ on \mathcal{S} , adopting $\mathbf{B}_H(\mathbf{q})$ as sight-restriction, and $\mathbf{B}_P(\mathbf{q})$ as evaluation function, with the induced preference relation.

³One might want to impose stronger variants of correctness. For instance the fact that if a player can see he will be moving again, then he will consider at least as much as he is considering now, from that history on: $\forall \mathbf{q} \in \mathbf{Q}$, if $\mathbf{q} = (h_0, h_1, \dots, h_k)$, $\mathbf{q}' = (h_1, \dots, h_k)$ and $t(h_0) = t(h_1)$ then $\mathbf{B}_H(\mathbf{q}') = \mathbf{B}_H(\mathbf{q})|_{h_0}$.

We can also impose that this fact is common knowledge among the players: $\forall \mathbf{q} \in \mathbf{Q}$, if $\mathbf{q} = (h_0, h_1, h_{i-1}, h_i, h_{i+1} \dots, h_k)$, $\mathbf{q}' = (h_0, h_1, h_{i-1}, h_{i+1} \dots, h_k)$ and $t(h_i) = t(h_{i-1})$ then $\mathbf{B}_H(\mathbf{q}') = \mathbf{B}_H(\mathbf{q})|_{h_i}$.

3.2 Analysing EMTGs

Example 2 has illustrated a natural notion of solution in an epistemic MTG, where each player calculates a best action in his or her sight restriction according to his or her evaluation criteria, recursively computing both the sight and the evaluation criteria of the other players. This is the idea behind the solution concept we propose for EMTGs.

Definition 6 (Nested Beliefs Solution) Let $\mathcal{S} = (S, \mathbf{B})$ be an EMTG and let $\mathbf{q} = (h_0, h_1, \dots, h_k)$ be a history-sequence. A strategy profile $\sigma \lceil_{\mathbf{B}(\mathbf{q})}$ is a Nested Beliefs Solution (NBS) of $\mathcal{S} \lceil_{\mathbf{B}(\mathbf{q})}$ if:

Base step For each $h' \in H \lceil_{\mathbf{B}_H(\mathbf{q})}$ that is a quasi-terminal history of $H \lceil_{\mathbf{B}_H(\mathbf{q})}$, we have that $h' \sigma \lceil_{\mathbf{B}_H(\mathbf{q}, h')}(h') \succeq_{t(h')}^{\mathbf{B}_P(\mathbf{q}, h')} h', \sigma' \lceil_{\mathbf{B}_H(\mathbf{q}, h')}(h')$ for any $\sigma' \lceil_{\mathbf{B}_H(\mathbf{q}, h')}$ that agrees with $\sigma \lceil_{\mathbf{B}(\mathbf{q})}$ up to h' .

Induction step For each $h' \in H \lceil_{\mathbf{B}_H(\mathbf{q})}$ that is neither terminal nor quasi-terminal in $H \lceil_{\mathbf{B}_H(\mathbf{q})}$, we have that

- $\sigma \lceil_{\mathbf{B}(\mathbf{q})}(h')$ agrees at h' with some Nested Beliefs Solution of $\mathcal{S} \lceil_{\mathbf{B}(\mathbf{q}, h')}$.
- If, for each $h' \in H \lceil_{\mathbf{B}_H(\mathbf{q})}$ that is neither terminal nor quasi-terminal in $H \lceil_{\mathbf{B}_H(\mathbf{q})}$, we have that $\sigma' \lceil_{\mathbf{B}(\mathbf{q})}(h')$ agrees at h' with some Nested Beliefs Solution of $\mathcal{S} \lceil_{\mathbf{B}(\mathbf{q}, h')}$ then the outcome z' generated by $\sigma' \lceil_{\mathbf{B}(\mathbf{q})}$ following h and the outcome z generated by $\sigma \lceil_{\mathbf{B}(\mathbf{q})}$ following h are such that $z \succeq_{t(h_k)}^{\mathbf{B}_P(\mathbf{q})} z'$.

We denote $\mathbf{NBS}(\mathcal{S} \lceil_{\mathbf{B}(\mathbf{q})})$ the set of NBS outcomes of $(\mathcal{S}, \mathbf{q})$. The composition of such outcomes yields our game solution.

Intuitively, a Nested Beliefs Solution of some game $\mathcal{S} \lceil_{\mathbf{B}(\mathbf{q})}$ is a best response to all Nested Belief Solutions at deeper level, e.g., of each $\mathcal{S} \lceil_{\mathbf{B}(\mathbf{q}, h')}$. Notice that because of the properties of sight functions the depth iteration is bound to reach a fixpoint.

Example 7 Let's go back to Figure 1 and compute the **NBS** at history A. We know there are four relevant histories sequences: (A), (A, C), (A, B) (A, B, C). To each of them we can associate the corresponding beliefs, as follows:

- $H \lceil_{\mathbf{B}_H(A)} = \{g, d, e, f, C, B, A\} = H \lceil_{h_0}$
- $H \lceil_{\mathbf{B}_H(A, B)} = \{e, C, f\}$
- $H \lceil_{\mathbf{B}_H(A, C)} = \{d\}$
- $H \lceil_{\mathbf{B}_H(A, B, C)} = \{e\}$

Let us know, for each histories sequence q specify the preference relation $\succeq_{t_k}^{\mathbf{B}_P(q)}$ (modulo reflexivity and transitivity), which is all we need to compute **NBS**.

- $\succeq_{Ann}^{\mathbf{B}_P(A)} = \{(d, g), (g, e), (e, f), (f, e)\}$
- $\succeq_{Bob}^{\mathbf{B}_P(A, B)} = \{(e, f), (f, d)\}$
- $\succeq_{Charles}^{\mathbf{B}_P(A, C)} = \{\}$

- $\succeq_{Charles}^{\mathbf{B}_P(A, B, C)} = \{\}$

Consider now the following strategy $\sigma \lceil_{\mathbf{B}(A)}$ ⁴

- $\sigma \lceil_{\mathbf{B}(A)}(A)_{Ann} = g$
- $\sigma \lceil_{\mathbf{B}(A)}(B)_{Bob} = C$
- $\sigma \lceil_{\mathbf{B}(A)}(C)_{Charles} = d$

Is σ a Nested Beliefs Solution of $\mathcal{S} \lceil_{\mathbf{B}(A)}$?

The condition at the base step is met by $\sigma \lceil_{\mathbf{B}(A)}(C)_{Charles} = d$.

Lets now look at $\sigma \lceil_{\mathbf{B}(A)}(B)_{Bob} = C$. Is C compatible with the best Nested Beliefs Solution of $\mathcal{S} \lceil_{\mathbf{B}(A, B)}$? We need first to compute all NBS of $\mathcal{S} \lceil_{\mathbf{B}(A, B)}$. Luckily there are not so many. Every such strategy must be of the form $\sigma' \lceil_{\mathbf{B}(A, B)}(C)_{Charles} = e$ and be the best among the strategies agreeing with NBS of $\mathcal{S} \lceil_{\mathbf{B}(A, B, C)}$ at C . So, given the preferences of B , be such that $\sigma' \lceil_{\mathbf{B}(A, B)}(B)_{Bob} = C$. This is indeed what σ does.

However notice that given the preference of A , σ is **not** behaving as a NBS at A , because Ann prefers d to g .

The strategy $\sigma^* \lceil_{\mathbf{B}(A)}$ only disagreeing with $\sigma \lceil_{\mathbf{B}(A)}$ at (A) , and being such that $\sigma^* \lceil_{\mathbf{B}(A)}(A)_{Ann} = B$, is a NBS of $\mathcal{S} \lceil_{\mathbf{B}(A)}$.

The composition of Nested Beliefs Solutions constitutes a rational outcome of the game.

Definition 8 (Sight-Compatible Epistemic Solution) Let $\mathcal{S} = (S, \mathbf{B})$ be an EMTG. A strategy profile σ is a Sight-Compatible Epistemic Solution (SCES) if at each $h \in H \setminus Z$, there exists a terminal history $z \in Z \lceil_{\mathbf{B}_H(h)}$ such that $h\sigma(h) \trianglelefteq z$ and $z \in \mathbf{NBS}(\mathcal{S} \lceil_{\mathbf{B}(h)})$.

We denote **SCES** the set of Sight-Compatible Epistemic Solutions of \mathcal{S} .

A SCES is the composition of best moves of players at each history. Each such move is a best response to what the current player believes other players will do and this belief is supported by all higher-order beliefs, compatible with the player's sight, about what the opponents can perceive and how they will evaluate it.

3.2.1 Computing rational solutions

Algorithm $Sol(\mathcal{S})$ below takes as input an EMTG and returns a path obtained by composing locally rational moves, compatible with players' higher-order beliefs about sights and evaluation criteria of their opponents. Algorithms 1 calls Algorithm 2, which in turn calls Algorithms 3 and 4. For technical convenience, we define VLP to be a dummy always dominated history.

The following theorem shows that every EMTG has a Sight-Compatible Epistemic Solution. Its proof consists in constructively building the desired strategy profile.

Theorem 9 (Existence Theorem) Let $\mathcal{S} = (S, \mathbf{B})$ be an EMTG. There exists a strategy profile σ that is a Sight-Compatible Epistemic Solution for $\mathcal{S} = (S, \mathbf{B})$.

⁴Slightly abusing notation, but unambiguously, we identify actions chosen by the strategy with the resulting histories.

Algorithm 1: Solution of \mathcal{S}

```

1 Sol( $\mathcal{S}$ )
  Input: An EMTG  $\mathcal{S} = (S, \mathbf{B})$ 
  Output: A terminal history  $h$  of  $\mathcal{S}$ 
2 begin
3    $h \leftarrow \varepsilon;$ 
4   while  $h \notin Z$  do
5      $h \leftarrow (h, BSBI(\mathcal{S}, h));$                                 /* NBS at  $h$  */
6   Return  $h;$ 

```

Algorithm 2: The current best move

```

1 BSBI( $\mathcal{S}, h$ )
  Input: A game  $\mathcal{S} = (S, \mathbf{B})$ , and a history  $h$ 
  Output: NBS move  $a$  at  $h$ 
2 begin
3   for each  $h' \in \mathbf{B}_H(h)$  and  $h' \neq h$  do
4      $Continuations[h'] \leftarrow NBS(S, (h, h'));$  /* Store NBS actions in an array, one for
       $each h'$  */
5   Return  $BB(\mathcal{S}, (h), Continuations);$ 

```

Proof Sketch: Let H be the set of histories in \mathcal{S} . For every $h \in H$ set $\sigma(h) := a$ for $a \leq h^*$ and h^* be the outcome returned by Algorithm 1 on input $\mathcal{S}|_{\mathbf{B}(h)}$. That the Algorithm returns a profile $o(\sigma)$ such that σ satisfies the conditions of Definition 6 at each history is a lengthy but relatively straightforward check, which we omit for space reasons. \square

Theorem 10 (Completeness Theorem) *Let $\mathcal{S} = (S, \mathbf{B})$ be a finite EMTG and let σ be a Sight-Compatible Epistemic Solution for $\mathcal{S} = (S, \mathbf{B})$. There exists an execution of Algorithm 1 returning $o(\sigma)$.*

Proof Sketch: Let $\mathcal{S} = (S, \mathbf{B})$ be a finite EMTG and let σ be a Sight-Compatible Epistemic Solution for $\mathcal{S} = (S, \mathbf{B})$. Now choose an execution of Algorithm 1 that is compatible with the action selection that, at each history sequence, is made by σ , which exists by construction. The finiteness assumption ensures termination. \square

The following observations illustrate the relation between SCES and the other two relevant solution concepts in the literature: SCBI (Grossi and Turrini, 2012) and classical BI (Osborne and Rubinstein, 1994). They specify precise conditions under which our solution concept collapses into these two.

Proposition 11 *Let $\mathcal{S} = (S, \mathbf{B})$ be an EMTG. If for any history-sequence $\mathbf{q} = (h_0, h_1, \dots, h_k)$ and any history $h' \in \mathbf{B}_H(\mathbf{q})$, $\mathbf{B}_H(\mathbf{q}, h') = \mathbf{B}_H(\mathbf{q})|_{h'}$, and $\mathbf{B}_P(\mathbf{q}) = B_P(h_0)$, then $\text{SCES}(\mathcal{S}) = \text{SCBI}(\mathcal{S})$.*

So, if the current player believes the following players' sights and evaluation criteria, together with their beliefs about other players' sights and evaluation criteria, are coherent with his', then SCES is equivalent to SCBI.

Algorithm 3: Beliefs of moves of following players

```

1 NBS( $\mathcal{S}$ ,  $\mathbf{q}$ )
  Input: A game  $\mathcal{S} = (S, \mathbf{B})$ , and a history sequence  $\mathbf{q} = (h_0, h_1, h_2, \dots, h_k)$ 
  Output: An action  $a$  following  $h_k$ 
2 begin
3   if  $h_k \in Z \cap \mathbf{B}_H(\mathbf{q})$  then
4     | Return  $\varepsilon$ ;
5   else
6     | for each  $h_{k+1} \in \mathbf{B}_H(\mathbf{q})$  and  $h_{k+1} \neq h_k$  do
7       |    $Continuations[h_{k+1}] \leftarrow NBS(S, (\mathbf{q}, h_{k+1}))$ ; /* Store NBS actions in an array,
      |   one for each  $h_{k+1}$  */
8   | Return  $BB(\mathcal{S}, \mathbf{q}, Continuations)$ ;

```

We know that the solution concept BI is a special case of SCBI, and therefore also of SCES.

Proposition 12 *Let $\mathcal{S} = (S, \mathbf{B})$ be an EMTG. If, for any history-sequence $\mathbf{q} = (h_0, h_1, \dots, h_k)$, we have that $\mathbf{B}_H(\mathbf{q}) = H|_{h_k}$ and that $\mathbf{B}_P(\mathbf{q}) = P_{t(h_k)}$ then $\mathbf{SCES}(\mathcal{S}) = \mathbf{BI}(\mathcal{S})$.*

The above result says that SCES coincides with standard backwards induction solution if, at each history, we have that higher-order beliefs about sight and evaluation criteria are *coherent* with the real subgame the current player faces and the preference relation the current player holds.

Despite the crucial presence of higher-order beliefs about sight-restricted games, we can show the following fairly surprising complexity result.

Proposition 13 *Given a finite EMTG \mathcal{S} , the problem of computing a SCES of \mathcal{S} is P-TIME complete.*

Proof sketch: For the upper bound, the key fact is that algorithm $Sol(S)$ runs in time $\mathcal{O}((n \log n)^2)$, with n being the cardinality of the set of histories of \mathcal{S} . This follows from the equations and facts below, where b and d are the largest number of branches and the depth of game tree respectively: 1). $T(sol(S)) = \mathcal{O}(T(BSBI * d))$; 2). $T(NBS) = \mathcal{O}(T(BSBI))$; 3). $T(BB) = \mathcal{O}(b * d)$; 4). Let $f(d) = T(BSBI)$, then $f(d) = \mathcal{O}(b * f(d-1) + b^2 * f(d-2) + b^3 * f(d-3) + \dots + b^{d-1} f(1) + T(BB)) = \mathcal{O}(d * 2^d * b^d)$; 5). $d \leq \log(n)$ and $2^d \leq b^d \leq n$. P-TIME hardness is a consequence of (Szymanik, 2013, Theorem 2), which shows that BI is P-TIME hard, and Proposition 12.

As a side remark, using a similar argument and Proposition 12 we are able to show that computing SCBI solutions is P-TIME complete.

4 Conclusions and potential developments

We have proposed a model for decision-making among resource-bounded players in extensive games, integrating an analytical perspective coming game theory with a procedural perspective coming from AI. In particular we have studied players with limited foresight which can

Algorithm 4: Best Branch

```
1 BB( $\mathcal{S}, \mathbf{q}, Continuations$ ) /* Compose chosen moves (in array Continuations), thus get all
   paths following  $h_k$  and choose a best move following  $h_k$  */
  Input: A game  $\mathcal{S}$ , a history sequence  $\mathbf{q}$ , an array Continuations
  Output: A best move following  $h_k$  determined by Continuations
2 begin
3    $bestpath \leftarrow VLP$ ;           /* VLP is a dominated history for all players */
4   for each  $(h_k, a) \in \mathbf{B}_H(\mathbf{q})$  do    /*  $a$  is any action following  $h_k$ , next we choose an
   optimal one in  $\mathbf{B}_H(\mathbf{q})$  */
5      $TP \leftarrow (h_k, a)$ ;
6     while Continuations[ $TP$ ] is defined in array Continuations do
7        $TP \leftarrow (TP, Continuations[TP])$ ;
8       if  $TP \succ_{t(h_k)}^{B_P(\mathbf{q})} bestpath$  then
9          $bestpath \leftarrow TP$ ;
10         $bestmove \leftarrow a$ ;
11   Return bestmove;
```

reason about their opponents, constructing beliefs about their limited abilities for calculation and evaluation, showing that our novel games have a well-behaved solution, generalising existing ones in the literature.

There are interesting modelling issues, as noted previously. Our game models strike a balance between simple trees as used for BI and more complex models as found in epistemic game theory (Perea, 2012). Here, what we left open is the relation between EMTGs and the Extensive Games with Awareness of (Halpern and Rêgo, 2006). We expect that the correspondence for GSSs of Theorem 3 in (Grossi and Turrini, 2012) can be lifted to EMTGs, using an iteration of the awareness functions Aw_i for players i to simulate the believed game at a history sequence. We stress, though, that the specific features of EMTGs give them an independent conceptual and technical interest. The emphasis on limited foresight (as opposed to perceiving a novel extensive game in (Halpern and Rêgo, 2006)) makes them a natural candidate for addressing Rubinstein's modelling challenge (Rubinstein, 2004), while still supporting an efficient algorithm to calculate the game equilibria.

Finally, our analysis raises several issues of logical definability and styles of reasoning. We believe that our solution concept is still definable in a computationally well-behaved logical language, a natural candidate being the fixed-point logic FOL(FP), shown in (van Benthem and Gheerbrant, 2010) to express backwards induction. What is new in our setting is that the reasoning underpinning our main theorems is a mixture of a backward induction style with a forward induction style (Perea, 2012; van Benthem, 2014), since we have to evaluate what players further down in the game tree are going to do according to players whose moves occurred earlier on in the game.

5 Acknowledgements

The author acknowledges the support of Imperial College London for the Junior Research Fellowship "Designing negotiation spaces for collective decision-making" (DoC- AI1048). He is besides extremely thankful to Johan van Benthem, Davide Grossi and Chanjuan Liu for their feedback on earlier versions of the paper.

References

Cameron Browne, Edward J. Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A Survey of Monte Carlo Tree Search Methods. *IEEE Trans. Comp. Intell. AI Games*, 4(1):1–43, 2012.

Jeroen Donkers, Jaap Van Den Herik, and Jos Uiterwijk. Probabilistic opponent-model search. *Information Sciences*, 135:3–4, 2001.

Prashant Doshi and Piotr J. Gmytrasiewicz. Monte carlo sampling methods for approximating interactive pomdps. *J. Artif. Intell. Res. (JAIR)*, 34:297–337, 2009.

Yossi Feinberg. Games with unawareness. *Stanford Graduate School of Business Paper No. 2122*, 2012.

Piotr J. Gmytrasiewicz and Prashant Doshi. A framework for sequential planning in multi-agent settings. *J. Artif. Intell. Res. (JAIR)*, 24:49–79, 2005.

Davide Grossi and Paolo Turrini. Short sight in extensive games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages 805–812, 2012.

Joseph Y. Halpern and Leandro Chaves Rêgo. Extensive games with possibly unaware players. In *AAMAS*, pages 744–751, 2006.

Joseph Y. Halpern. Beyond nash equilibrium: Solution concepts for the 21st century. *CoRR*, abs/0806.2139, 2008.

Aviad Heifetz, Martin Meier, and Burkhard C. Schipper. Interactive unawareness. *J. Economic Theory*, 130(1):78–94, 2006.

Aviad Heifetz, Martin Meier, and Burkhard C. Schipper. Dynamic unawareness and rationalizable behavior. *Games and Economic Behavior*, 81:50–68, 2013.

Aviad Heifetz, Martin Meier, and Burkhard C. Schipper. Unawareness, beliefs, and speculative trade. *Games and Economic Behavior*, 77(1):100–121, 2013.

Martin J Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

Andrés Perea. *Epistemic Game Theory: Reasoning and Choice*. Cambridge University Press, 2012.

Ariel Rubinstein. *Modelling Bounded Rationality*. MIT Press, 2004.

Stuart Russel and Peter Norvig. *Artificial Intelligence: A Modern Approach (3rd edition)*. Pearson, 2012.

Stuart J. Russell and Eric Wefald. Principles of metareasoning. *Artif. Intell.*, 49(1-3):361–395, 1991.

Frederik Schadd, Sander Bakkes, and Pieter Spronck. Opponent modeling in real-time strategy games. In *Proceedings of the 8th International Conference on Intelligent Games and Simulation (GAMEON'2007)*, pages 61–70, 2007.

Ekhlas Sonu and Prashant Doshi. Scalable solutions of interactive pomdps using generalized and bounded policy iteration. *Autonomous Agents and Multi-Agent Systems*, 29(3):455–494, 2015.

Jakub Szymanik. Backward Induction Is PTIME-complete. In Davide Grossi, Olivier Roy, and Huaxin Huang, editors, *Logic, Rationality, and Interaction - 4th International Workshop, LORI 2013, Hangzhou, China, October 9-12, 2013, Proceedings*, volume 8196 of *Lecture Notes in Computer Science*, pages 352–356. Springer, 2013.

Johan van Benthem and Amélie Gheerbrant. Game solution, epistemic dynamics and fixed-point logics. *Fundam. Inform.*, 100(1-4):19–41, 2010.

Johan van Benthem, Eric Pacuit, and Olivier Roy. Toward a theory of play: A logical perspective on games and interaction. *Games*, 2(1):52–86, 2011.

Johan van Benthem. *Logic in Games*. MIT Press, 2014.